

Exhibits A through F filed under seal

Exhibit G



Let's Take the Con Out of Econometrics

Edward E. Leamer

The American Economic Review, Volume 73, Issue 1 (Mar., 1983), 31-43.

Stable URL:

<http://links.jstor.org/sici?sici=0002-8282%28198303%2973%3A1%3C31%3ALTTCOO%3E2.0.CO%3B2-R>

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

The American Economic Review is published by American Economic Association. Please contact the publisher for further permissions regarding the use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/aea.html>.

The American Economic Review
©1983 American Economic Association

JSTOR and the JSTOR logo are trademarks of JSTOR, and are Registered in the U.S. Patent and Trademark Office. For more information on JSTOR contact jstor-info@umich.edu.

©2002 JSTOR

Let's Take the Con out of Econometrics

By EDWARD E. LEAMER*

Econometricians would like to project the image of agricultural experimenters who divide a farm into a set of smaller plots of land and who select randomly the level of fertilizer to be used on each plot. If some plots are assigned a certain amount of fertilizer while others are assigned none, then the difference between the mean yield of the fertilized plots and the mean yield of the unfertilized plots is a measure of the effect of fertilizer on agricultural yields. The econometrician's humble job is only to determine if that difference is large enough to suggest a real effect of fertilizer, or is so small that it is more likely due to random variation.

This image of the applied econometrician's art is grossly misleading. I would like to suggest a more accurate one. The applied econometrician is like a farmer who notices that the yield is somewhat higher under trees where birds roost, and he uses this as evidence that bird droppings increase yields. However, when he presents this finding at the annual meeting of the American Ecological Association, another farmer in the audience objects that he used the same data but came up with the conclusion that moderate amounts of shade increase yields. A bright chap in the back of the room then observes that these two hypotheses are indistinguishable, given the available data. He mentions the phrase "identification problem," which, though no one knows quite what he means, is said with such authority that it is totally convincing. The meeting reconvenes in the halls and in the bars, with heated discussion whether this is the kind of work that merits promotion from Associate to Full Farmer; the Luminists strongly opposed to promotion and the Aviophiles equally strong in favor.

One should not jump to the conclusion that there is necessarily a substantive difference between drawing inferences from experimental as opposed to nonexperimental data. The images I have drawn are deliberately prejudicial. First, we had the experimental scientist with hair neatly combed, wide eyes peering out of horn-rimmed glasses, a white coat, and an electronic calculator for generating the random assignment of fertilizer treatment to plots of land. This seems to contrast sharply with the nonexperimental farmer with overalls, unkempt hair, and bird droppings on his boots. Another image, drawn by Orcutt, is even more damaging: "Doing econometrics is like trying to learn the laws of electricity by playing the radio." However, we need not now submit to the tyranny of images, as many of us have in the past.

I. Is Randomization Essential?

What is the real difference between these two settings? Randomization seems to be the answer. In the experimental setting, the fertilizer treatment is "randomly" assigned to plots of land, whereas in the other case nature did the assignment. Now it is the tyranny of words that we must resist. "Random" does not mean adequately mixed in *every* sample. It only means that on the average, the fertilizer treatments are adequately mixed. Randomization implies that the least squares estimator is "unbiased," but that definitely does not mean that for each sample the estimate is correct. Sometimes the estimate is too high, sometimes too low. I am reminded of the lawyer who remarked that "when I was a young man I lost many cases that I should have won, but when I grew older I won many that I should have lost, so on the average justice was done."

In particular, it is possible for the randomized assignment to lead to exactly the same allocation as the nonrandom assignment,

*Professor of economics, University of California-Los Angeles. This paper was a public lecture presented at the University of Toronto, January 1982. I acknowledge partial support by NSF grant SOC78-09479.

namely, with treated plots of land all being under trees and with nontreated plots of land all being away from trees. I submit that, if this is the outcome of the randomization, then the randomized experiment and the nonrandomized experiment are exactly the same. Many econometricians would insist that there is a difference, because the randomized experiment generates "unbiased" estimates. But all this means is that, if this particular experiment yields a gross overestimate, some other experiment yields a gross underestimate.

Randomization thus does not assure that each and every experiment is "adequately mixed," but randomization does make "adequate mixing" probable. In order to make clear what I believe to be the true value of randomization, let me refer to the model

$$(1) \quad Y_i = \alpha + \beta F_i + \gamma L_i + U_i,$$

where Y_i is the yield of plot i ; F_i is the fertilizer assigned to plot i ; L_i is the light falling on plot i ; U_i is the unspecified influence on the yield of plot i , and where β , the fertilizer effect, is the object of the inferential exercise. We may suppose to begin the argument that the light level is expensive to measure and that it is decided to base an estimate of β initially only on measurement of Y_i and F_i . We may assume also that the natural experiment produces values for F_i , L_i , and U_i with expected values $E(U_i|F_i) = 0$ and $E(L_i|F_i) = r_0 + r_1 F_i$. In the more familiar parlance, it is assumed that the fertilizer level and the residual effects are uncorrelated, but the fertilizer level and the light level are possibly correlated. As every beginning econometrics student knows, if you omit from a model a variable which is correlated with included variables, bad things happen. These bad things are revealed to the econometrician by computing the conditional mean of Y given F but not L :

$$\begin{aligned} (2) \quad E(Y|F) &= \alpha + \beta F + \gamma E(L|F) \\ &= \alpha + \beta F + \gamma(r_0 + r_1 F) \\ &\equiv (\alpha + \alpha^*) + (\beta + \beta^*)F, \end{aligned}$$

where $\alpha^* = \gamma r_0$ and $\beta^* = \gamma r_1$. The linear regression of Y on F provides estimates of the parameters of the conditional distribution of Y given F , and in this case the regression coefficients are estimates not of α and β , but rather of $\alpha + \alpha^*$ and $\beta + \beta^*$. The parameters α^* and β^* measure the bias in the least squares estimates. This bias could be due to left-out variables, or to measurement errors in F , or to simultaneity.

When observing a nonexperiment, the bias parameters α^* and β^* can be thought to be small, but they cannot sensibly be treated as exact zeroes. The notion that the bias parameters are small can be captured by the assumption that α^* and β^* are drawn from a normal distribution with zero means and covariance matrix M . The model can then be written as $Y = \alpha + \beta F + \varepsilon$, where ε is the sum of three random variables: $U + \alpha^* + \beta^* F$. Because the error term ε is not spherical, the proper way to estimate α and β is generalized least squares. My 1974 article demonstrates that if (a, b) represent the least squares estimates of (α, β) , then the generalized least squares estimates (\hat{a}, \hat{b}) are also equal to (a, b) :

$$(3) \quad \begin{pmatrix} \hat{a} \\ \hat{b} \end{pmatrix} = \begin{pmatrix} a \\ b \end{pmatrix},$$

and if S represents the sample covariance matrix for the least squares estimates, then the sample covariance matrix for (\hat{a}, \hat{b}) is

$$(4) \quad \text{Var}(\hat{a}, \hat{b}) = S + M,$$

where M is the covariance matrix of (α^*, β^*) .

The meaning of equation (3) is that unless one knows the direction of the bias, the possibility of bias does not call for any adjustment to the estimates. The possibility of bias does require an adjustment to the covariance matrix (4). The uncertainty is composed of two parts: the usual sampling uncertainty S plus the misspecification uncertainty M . As sample size grows, the sampling uncertainty S ever decreases, but the misspecification uncertainty M remains ever constant. The misspecification matrix M that we must add to the least squares variance

matrix is just the (prior) variance of the bias coefficients (α^* , β^*). If this variance matrix is small, the least squares bias is likely to be small. If M is large, it is correspondingly probable that (α^* , β^*) is large.

It would be a remarkable bootstrap if we could determine the extent of the misspecification from the data. The data in fact contain no information about the size of the bias, a point which is revealed by studying the likelihood function. The misspecification matrix M is therefore a pure prior concept. One must decide independent of the data how good the nonexperiment is.

The formal difference between a randomized experiment and a natural experiment is measured by the matrix M . If the treatment is randomized, the bias parameters (α^* , β^*) are exactly zero, or, equivalently, the matrix M is a zero matrix. If M is zero, the least squares estimates are consistent. If M is not zero, as in the natural experiment, there remains a fixed amount of specification uncertainty, independent of sample size.

There is therefore a sharp difference between inference from randomized experiments and inference from natural experiments. This seems to draw a sharp distinction between economics where randomized experiments are rare and "science" where experiments are routinely done. But the fact of the matter is that no one has ever designed an experiment that is free of bias, and no one can. As it turns out, the technician who was assigning fertilizer levels to plots of land, took his calculator into the fields, and when he was out in the sun, the calculator got heated up and generated large "random" numbers, which the technician took to mean no fertilizer; and when he stood under the shade of the trees, his cool calculator produced small numbers, and these plots received fertilizer.

You may object that this story is rather fanciful, but I need only make you think it is possible, to force you to set $M \neq 0$. Or if you think a computer can really produce random numbers (calculated by a mathematical formula and therefore perfectly predictable!), I will bring up mismeasurement of the fertilizer level, or human error in carrying out the computer instructions. Thus, the attempt to

randomize and the attempt to measure accurately ensures that M is small, but not zero, and the difference between scientific experiments and natural experiments is difference in degree, but not in kind. Admittedly however, the misspecification uncertainty in many experimental settings may be so small that it is well approximated by zero. This can very rarely be said in nonexperimental settings.

Examples may be ultimately convincing. There is a great deal of empirical knowledge in the science of astronomy, yet there are no experiments. Medical knowledge is another good example. I was struck by a headline in the January 5, 1982 *New York Times*: "Life Saving Benefits of Low-Cholesterol Diet Affirmed in *Rigorous* Study." The article describes a randomized experiment with a control group and a treated group. "*Rigorous*" is therefore interpreted as "randomized." As a matter of fact, there was a great deal of evidence suggesting a link between heart disease and diet before any experiments were performed on humans. There were cross-cultural comparisons and there were animal studies. Actually, the only reason for performing the randomized experiment was that someone believed there was pretty clear non-experimental evidence to begin with. The nonexperimental evidence was, of course, inconclusive, which in my language means that the misspecification uncertainty M remained uncomfortably large. The fact that the Japanese have both less incidence of heart disease and also diets lower in cholesterol compared to Americans is not convincing evidence, because there are so many other factors that remain unaccounted for. The fact that pigs on a high cholesterol diet develop occluded arteries is also not convincing, because the similarity in physiology in pigs and humans can be questioned.

When the sampling uncertainty S gets small compared to the misspecification uncertainty M , it is time to look for other forms of evidence, experiments or nonexperiments. Suppose I am interested in measuring the width of a coin, and I provide rulers to a room of volunteers. After each volunteer has reported a measurement, I compute the mean and standard deviation, and I conclude that

the coin has width 1.325 millimeters with a standard error of .013. Since this amount of uncertainty is not to my liking, I propose to find three other rooms full of volunteers, thereby multiplying the sample size by four, and dividing the standard error in half. That is a silly way to get a more accurate measurement, because I have already reached the point where the sampling uncertainty S is very small compared with the misspecification uncertainty M . If I want to increase the true accuracy of my estimate, it is time for me to consider using a micrometer. So too in the case of diet and heart disease. Medical researchers had more or less exhausted the vein of nonexperimental evidence, and it became time to switch to the more expensive but richer vein of experimental evidence.

In economics, too, we are switching to experimental evidence. There are the laboratory experiments of Charles Plott and Vernon Smith (1978) and Smith (1980), and there are the field experiments such as the Seattle/Denver income maintenance experiment. Another way to limit the misspecification error M is to gather different kinds of nonexperiments. Formally speaking, we will say that experiment 1 is qualitatively different from experiment 2 if the bias parameters (α_1^*, β_1^*) are distributed independently of the bias parameters (α_2^*, β_2^*) . In that event, simple averaging of the data from the two experiments yields average bias parameters $(\alpha_1^* + \alpha_2^*, \beta_1^* + \beta_2^*)/2$ with misspecification variance matrix $M/2$, half as large as the (common) individual variances. Milton Friedman's study of the permanent income hypothesis is the best example of this that I know. Other examples are hard to come by. I believe we need to put much more effort into identifying qualitatively different and convincing kinds of evidence.

Parenthetically, I note that traditional econometric theory, which does not admit experimental bias, as a consequence also admits no "hard core" propositions. Demand curves can be shown to be positively sloped. Utility can be shown not to be maximized. Econometric evidence of a positively sloped demand curve would, as a matter of fact, be routinely explained in terms of simultaneity bias. If utility seems not to have been maxi-

mized, it is only that the econometrician has misspecified the utility function. The misspecification matrix M thus forms Imre Lakatos' "protective belt" which protects certain hard core propositions from falsification.

II. Is Control Essential?

The experimental scientist who notices that the fertilizer treatment is correlated with the light level can correct his experimental design. He can control the light level, or he can allocate the fertilizer treatment in such a way that the fertilizer level and the light level are not perfectly correlated.

The nonexperimental scientist by definition cannot control the levels of extraneous influences such as light. But he can control for the variable light level by including light in the estimating equation. Provided nature does not select values for light and values for fertilizer levels that are perfectly correlated, the effect of fertilizer on yields can be estimated with a multiple regression. The collinearity in naturally selected treatment variables may mean that the data evidence is weak, but it does not invalidate in any way the usual least squares estimates. Here, again, there is no essential difference between experimental and nonexperimental inference.

III. Are the Degrees of Freedom Inadequate with Nonexperimental Data?

As a substitute for experimental control, the nonexperimental researcher is obligated to include in the regression equation all variables that might have an important effect. The NBER data banks contain time-series data on 2,000 macroeconomic variables. A model explaining gross national product in terms of all these variables would face a severe degrees-of-freedom deficit since the number of annual observations is less than thirty. Though the number of observations of any phenomenon is clearly limited, the number of explanatory variables is logically unlimited. If a polynomial could have a degree as high as k , it would usually be admitted that the degree could be $k+1$ as well. A theory that allows k lagged explanatory vari-




FIGURE 1. HYPOTHETICAL DATA AND THREE ESTIMATED QUADRATIC FUNCTIONS

ables would ordinarily allow $k+1$. If the level of money might affect GNP , then why not the number of presidential sneezes, or the size of the polar ice cap?

The number of explanatory variables is unlimited in a nonexperimental setting, but it is also unlimited in an experimental setting. Consider again the fertilizer example in which the farmer randomly decides either to apply F_1 pounds of fertilizer per acre or zero pounds, and obtains the data illustrated in Figure 1. These data admit the inference that fertilizer level F_1 produces higher yields than no fertilizer. But the farmer is interested in selecting the fertilizer level that maximizes profits. If it is hypothesized that yield is a linear function of the fertilizer intensity $Y = \alpha + \beta F + U$, then profits are

$$\text{Profits} = pA(\alpha + \beta F + U) - p_F AF,$$

where A is total acreage, p is the product price, and p_F is the price per pound of fertilizer. This profit function is linear in F with slope $A(\beta p - p_F)$. The farmer maximizes profits therefore by using no fertilizer if the price of fertilizer is high, $\beta p < p_F$, and using an unlimited amount of fertilizer if the price is low, $\beta p > p_F$. It is to be expected that you will find this answer unacceptable for one of




FIGURE 2. HYPOTHETICAL DATA AND ESTIMATED QUADRATIC FUNCTION

several reasons:

1) When the farmer tries to buy an unlimited amount of fertilizer, he will drive up its price, and the problem should be reformulated to make p_F a function of F .

2) Uncertainty in the fertilizer effect β causes uncertainty in profits, $\text{Variance}(\text{profits}) = p^2 A^2 F^2 \text{Var}(\beta)$, and risk aversion will limit the level of fertilizer applied.

3) The yield function is nonlinear.

Economic theorists doubtless find reasons 1) and 2) compelling, but I suspect that the real reason farmers don't use huge amounts of fertilizer is that the marginal increase in the yield eventually decreases. Plants don't grow in fertilizer alone.

So let us suppose that yield is a quadratic function of fertilizer intensity, $Y = \alpha + \beta_1 F + \beta_2 F^2 + U$, and suppose we have only the data illustrated in Figure 1. Unfortunately, there are an infinite number of quadratic functions all of which fit the data equally well, three of which are drawn. If there were no other information available, we could conclude only that the yield is higher at F_1 than at zero. Formally speaking, there is an identification problem, which can be solved by altering the experimental design. The yield must be observed at a third point, as in Figure 2, where I have drawn the least squares estimated quadratic function and have indicated the fertilizer intensity F_m that maximizes the yield. I expect that most people would question whether these data admit the




FIGURE 3. HYPOTHETICAL DATA AND THREE ESTIMATED FUNCTIONS

inference that the yield is maximized at F_m . Actually, after inspection of this figure, I don't think anything can be inferred except that the yield at F_2 is higher than at F_1 , which in turn is higher than at zero. Thus I don't believe the function is quadratic. If it is allowed to be a cubic then again there is an identification problem.

This kind of logic can be extended indefinitely. One can always find a set of observations that will make the inferences implied by a polynomial of degree p seem silly. This is true regardless of the degree p . Thus no model with a finite number of parameters is actually believed, whether the data are experimental or nonexperimental.

IV. Do We Need Prior Information?

A model with an infinite number of parameters will allow inference from a finite data set only if there is some prior information that effectively constrains the ranges of the parameters. Figure 3 depicts another hypothetical sequence of observations and three estimated relationships between yield and fertilizer. I believe the solid line *A* is a better representation of the relationship than either of the other two. The piecewise linear form *B* fits the data better, but I think this peculiar meandering function is highly unlikely on an a priori basis. Though *B* and *C* fit the data equally well, I believe that *B* is much more

likely than *C*. What I am revealing is the a priori opinion that the function is likely to be smooth and single peaked.

What should now be clear is that data alone cannot reveal the relationship between yield and fertilizer intensity. Data can reveal the yield at sampled values of fertilizer intensities, but in order to interpolate between these sampled values, we must resort to subjective prior information.

Economists have inherited from the physical sciences the myth that scientific inference is objective, and free of personal prejudice. This is utter nonsense. All knowledge is human belief; more accurately, human opinion. What often happens in the physical sciences is that there is a high degree of conformity of opinion. When this occurs, the opinion held by most is asserted to be an objective fact, and those who doubt it are labelled "nuts." But history is replete with examples of opinions losing majority status, with once-objective "truths" shrinking into the dark corners of social intercourse. To give a trivial example, coming now from California I am unsure whether fat ties or thin ties are aesthetically more pleasing.

The false idol of objectivity has done great damage to economic science. Theoretical econometricians have interpreted scientific objectivity to mean that an economist must identify exactly the variables in the model, the functional form, and the distribution of the errors. Given these assumptions, and given a data set, the econometric method produces an objective inference from a data set, unencumbered by the subjective opinions of the researcher.

This advice could be treated as ludicrous, except that it fills all the econometric textbooks. Fortunately, it is ignored by applied econometricians. The econometric art as it is practiced at the computer terminal involves fitting many, perhaps thousands, of statistical models. One or several that the researcher finds pleasing are selected for reporting purposes. This searching for a model is often well intentioned, but there can be no doubt that such a specification search invalidates the traditional theories of inference. The concepts of unbiasedness, consistency, efficiency, maximum-likelihood estimation,

in fact, all the concepts of traditional theory, utterly lose their meaning by the time an applied researcher pulls from the bramble of computer output the one thorn of a model he likes best, the one he chooses to portray as a rose. The consuming public is hardly fooled by this chicanery. The econometrician's shabby art is humorously and disparagingly labelled "data mining," "fishing," "grubbing," "number crunching." A joke evokes the Inquisition: "If you torture the data long enough, Nature will confess" (Coase). Another suggests methodological fickleness: "Econometricians, like artists, tend to fall in love with their models" (wag unknown). Or how about: "There are two things you are better off not watching in the making: sausages and econometric estimates."

This is a sad and decidedly unscientific state of affairs we find ourselves in. Hardly anyone takes data analyses seriously. Or perhaps more accurately, hardly anyone takes anyone else's data analyses seriously. Like elaborately plumed birds who have long since lost the ability to procreate but not the desire, we preen and strut and display our *t*-values.

If we want to make progress, the first step we must take is to discard the counterproductive goal of objective inference. The dictionary defines an inference as a logical conclusion based on a set of facts. The "facts" used for statistical inference about θ are first the data, symbolized by x , second a conditional probability density, known as a sampling distribution, $f(x|\theta)$, and, third, explicitly for a Bayesian and implicitly for "all others," a marginal or prior probability density function $f(\theta)$. Because both the sampling distribution and the prior distribution are actually *opinions* and not *facts*, a statistical inference is and must forever remain an *opinion*.

What is a fact? A fact is merely an opinion held by all, or at least held by a set of people you regard to be a close approximation to all.¹ For some that set includes only one

person. I myself have the opinion that Andrew Jackson was the sixteenth president of the United States. If many of my friends agree, I may take it to be a fact. Actually, I am most likely to regard it to be a fact if the authors of one or more books say it is so.

The difference between a fact and an opinion for purposes of decision making and inference is that when I use opinions, I get uncomfortable. I am not too uncomfortable with the opinion that error terms are normally distributed because most econometricians make use of that assumption. This observation has deluded me into thinking that the opinion that error terms are normal may be a fact, when I know deep inside that normal distributions are actually used only for convenience. In contrast, I am *quite* uncomfortable using a prior distribution, mostly I suspect because hardly anyone uses them. If convenient prior distributions were used as often as convenient sampling distributions, I suspect that I could be as easily deluded into thinking that prior distributions are facts as I have been into thinking that sampling distributions are facts.

To emphasize this hierarchy of statements, I display them in order: truths; facts; opinions; conventions. Note that I have added to the top of the order, the category truths. This will appeal to those of you who feel compelled to believe in such things. At the bottom are conventions. In practice, it may be difficult to distinguish a fact from a convention, but when facts are clearly unavailable, we must strongly resist the deceit or delusion that conventions can represent.

What troubles me about using opinions is their whimsical nature. Some mornings when I arise, I have the opinion that Raisin Bran is better than eggs. By the time I get to the kitchen, I may well decide on eggs, or oatmeal. I usually do recall that the sixteenth president distinguished himself. Sometimes I think he was Jackson; often I think he was Lincoln.

A data analysis is similar. Sometimes I take the error terms to be correlated, sometimes uncorrelated; sometimes normal and sometimes nonnormal; sometimes I include observations from the decade of the fifties, sometimes I exclude them; sometimes the

¹This notion of "truth by consensus" is espoused by Thomas Kuhn (1962) and Michael Polanyi (1964). Oscar Wilde agrees by dissent: "A truth ceases to be true when more than one person believes it."

equation is linear and sometimes nonlinear; sometimes I control for variable z , sometimes I don't. Does it depend on what I had for breakfast?

As I see it, the fundamental problem facing econometrics is how adequately to control the whimsical character of inference, how sensibly to base inferences on opinions when facts are unavailable. At least a partial solution to this problem has already been formed by practicing econometricians. A common reporting style is to record the inferences implied by alternative sets of opinions. It is not unusual to find tables that show how an inference changes as variables are added to or deleted from the equation. This kind of sensitivity analysis reports special features of the mapping from the space of assumptions to the space of inferences. The defect of this style is that the coverage of assumptions is infinitesimal, in fact a zero volume set in the space of assumptions. What is needed instead is a more complete, but still economical way to report the mapping of assumptions into inferences. What I propose to do is to develop a correspondence between regions in the assumption space and regions in the inference space. I will report that all assumptions in a certain set lead to essentially the same inference. Or I will report that there are assumptions within the set under consideration that lead to radically different inferences. In the latter case, I will suspend inference and decision, or I will work harder to narrow the set of assumptions.

Thus what I am asserting is that the choice of a particular sampling distribution, or a particular prior distribution, is inherently whimsical. But statements such as "The sampling distribution is symmetric and unimodal" and "My prior is located at the origin" are not necessarily whimsical, and in certain circumstances do not make me uncomfortable.

To put this somewhat differently, an inference is not believable if it is fragile, if it can be reversed by minor changes in assumptions. As consumers of research, we correctly reserve judgment on an inference until it stands up to a study of fragility, usually by other researchers advocating opposite opinions. It is, however, much more efficient for

individual researchers to perform their own sensitivity analyses, and we ought to be demanding much more complete and more honest reporting of the fragility of claimed inferences.

The job of a researcher is then to report economically and informatively the mapping from assumptions into inferences. In a slogan, "The mapping is the message." The mapping does not depend on opinions (assumptions), but reporting the mapping economically and informatively does. A researcher has to decide which assumptions or which sets of alternative assumptions are worth reporting. A researcher is therefore forced either to anticipate the opinions of his consuming public, or to recommend his own opinions. It is actually a good idea to do both, and a serious defect of current practice is that it concentrates excessively on convincing one's self and, as a consequence, fails to convince the general professional audience.

The whimsical character of econometric inference has been partially controlled in the past by an incomplete sensitivity analysis. It has also been controlled by the use of conventions. The normal distribution is now so common that there is nothing at all whimsical in its use. In some areas of study, the list of variables is partially conventional, often based on whatever list the first researcher happened to select. Even conventional prior distributions have been proposed and are used with nonnegligible frequency. I am referring to Robert Shiller's (1973) smoothness prior for distributed lag analysis and to Arthur Hoerl and Robert Kennard's (1970) ridge regression prior. It used to aggravate me that these methods seem to find public favor whereas overt and complete Bayesian methods such as my own proposals (1972) for distributed lag priors are generally ignored. However, there is a very good reason for this: the attempt to form a prior distribution from scratch involves an untold number of partly arbitrary decisions. The public is rightfully resistant to the whimsical inferences which result, but at the same time is receptive to the use of priors in ways that control the whimsy. Though the use of conventions does control the whimsy, it can do so at the cost of relevance. Inferences based

on Hoerl and Kennard's conventional "ridge regression" prior are usually irrelevant, because it is rarely sensible to take the prior to be spherical and located at the origin, and because a closer approximation to prior belief can be suspected to lead to substantially different inferences. In contrast, the conventional assumption of normality at least uses a distribution which usually cannot be ruled out altogether. Still, we may properly demand a demonstration that the inferences are insensitive to this distributional assumption.

A. *The Horizon Problem: Sherlock Holmes Inference*

Conventions are not to be ruled out altogether, however. One can go mad trying to report completely the mapping from assumptions into inferences since the space of assumptions is infinite dimensional. A formal statistical analysis therefore has to be done within the limits of a reasonable horizon. An informed convention can usefully limit this horizon. If it turned out that sensible neighborhoods of distributions around the normal distribution 99 times out of 100 produced the same inference, then we could all agree that there are other more important things to worry about, and we may properly adopt the convention of normality. The consistency of least squares estimates under wide sets of assumptions is used improperly as support for this convention, since the inferences from a given finite sample may nonetheless be quite sensitive to the normality assumption.²

The truly sharp distinction between inference from experimental and inference from nonexperimental data is that experimental inference sensibly admits a conventional horizon in a critical dimension, namely the choice of explanatory variables. If fertilizer is randomly assigned to plots of land, it is conventional to restrict attention to the relationship between yield and fertilizer, and

to proceed as if the model were perfectly specified, which in my notation means that the misspecification matrix M is the zero matrix. There is only a small risk that when you present your findings, someone will object that fertilizer and light level are correlated, and there is an even smaller risk that the conventional zero value for M will lead to inappropriate inferences. In contrast, it would be foolhardy to adopt such a limited horizon with nonexperimental data. But if you decide to include light level in your horizon, then why not rainfall; and if rainfall, then why not temperature; and if temperature, then why not soil depth, and if soil depth, then why not the soil grade; ad infinitum. Though this list is never ending, it can be made so long that a nonexperimental researcher can feel as comfortable as an experimental researcher that the risk of having his findings upset by an extension of the horizon is very low. The exact point where the list is terminated must be whimsical, but the inferences can be expected not to be sensitive to the termination point if the horizon is wide enough.

Still, the horizon within which we all do our statistical analyses has to be ultimately troublesome, since there is no formal way to know what inferential monsters lurk beyond our immediate field of vision. "Diagnostic" tests with explicit alternative hypotheses such as the Durbin-Watson test for first-order autocorrelation do not truly ask if the horizon should be extended, since first-order autocorrelation is explicitly identified and clearly in our field of vision. Diagnostic tests such as goodness-of-fit tests, without explicit alternative hypotheses, are useless since, if the sample size is large enough, any maintained hypothesis will be rejected (for example, no observed distribution is exactly normal). Such tests therefore degenerate into elaborate rituals for measuring the effective sample size.

The only way I know to ask the question whether the horizon is wide enough is to study the anomalies of the data. In the words of the physiologist, C. Bernard:

A great surgeon performs operations for stones by a single method; later he

²In particular, least squares estimates are completely sensitive to the independence assumption, since by choice of sample covariance matrix a generalized least squares estimate can be made to assume any value whatsoever (see my 1981 paper).

makes a statistical summary of deaths and recoveries, and he concludes from these statistics that the mortality law for this operation is two out of five. Well, I say that this ratio means literally nothing scientifically, and gives no certainty in performing the next operation. What really should be done, instead of gathering facts empirically, is to study them more accurately, each in its special determinism...by statistics, we get a conjecture of greater or less probability about a given case, but never any certainty, never any absolute determinism...only basing itself on experimental determinism can medicine become a true science.

[1927, pp. 137–38]

A study of the anomalies of the data is what I have called “Sherlock Holmes” inference, since Holmes turns statistical inference on its head: “It is a capital mistake to theorize before you have all the evidence. It biases the judgements.” Statistical theory counsels us to begin with an elicitation of opinions about the sampling process and its parameters; the theory, in other words. After that, data may be studied in a purely mechanical way. Holmes warns that this biases the judgements, meaning that a theory constructed before seeing the facts can be disastrously inappropriate and psychologically difficult to discard. But if theories are constructed after having studied the data, it is difficult to establish by how much, if at all, the data favor the data-instigated hypothesis. For example, suppose I think that a certain coefficient ought to be positive, and my reaction to the anomalous result of a negative estimate is to find another variable to include in the equation so that the estimate is positive. Have I found evidence that the coefficient is positive? It would seem that we should require evidence that is more convincing than the traditional standard. I have proposed a method for discounting such evidence (1974). Initially, when you regress yield on fertilizer as in equation (2), you are required to assess a prior distribution for the experimental bias parameter β^* ; that is, you must select the misspecification matrix M . Then, when the least squares estimate of β

turns out to be negative, and you decide to include in the equation the light level as well as the fertilizer level, you are obligated to form a prior for the light coefficient γ consistent with the prior for β^* , given that $\beta^* = \gamma r_1$, where r_1 is the regression coefficient of light on fertilizer.³

This method for discounting the output of exploratory data analysis requires a discipline that is lacking even in its author. It is consequently important that we reduce the risk of Holmesian discoveries by extending the horizon reasonably far. The degree of a polynomial or the order of a distributed lag need not be data instigated, since the horizon is easily extended to include high degrees and high orders. It is similarly wise to ask yourself before examining the data what you would do if the estimate of your favorite coefficient had the wrong sign. If that makes you think of a specific left-out variable, it is better to include it from the beginning.

Though it is wise to select a wide horizon to reduce the risk of Holmesian discoveries, it is mistaken then to analyze a data set as if the horizon were wide enough. Within the limits of a horizon, no revolutionary inference can be made, since all possible inferences are predicted in advance (admittedly, some with low probabilities). Within the horizon, inference and decision can be turned over completely to a computer. But the great human revolutionary discoveries are made when the horizon is extended for reasons that cannot be predicted in advance and cannot be computerized. If you wish to make such discoveries, you will have to poke at the horizon, and poke again.

V. An Example

This rhetoric is understandably tiring. Methodology, like sex, is better demonstrated than discussed, though often better anticipated than experienced. Accordingly, let me give you an example of what all this

³In a randomized experiment with $r_1 = 0$, the constraint $\beta^* = \gamma r_1$ is irrelevant, and you are free to play these exploratory games without penalty. This is a very critical difference between randomized experiments and nonrandomized nonexperiments.

ranting and raving is about. I trust you will find it even better in the experience than in the anticipation. A problem of considerable policy importance is whether or not to have capital punishment. If capital punishment had no deterrent value, most of us would prefer not to impose such an irreversible punishment, though, for a significant minority, the pure joy of vengeance is reason enough. The deterrent value of capital punishment is, of course, an empirical issue. The unresolved debate over its effectiveness began when evolution was judging the survival value of the vengeance gene. Nature was unable to make a decisive judgment. Possibly econometricians can.

In Table 1, you will find a list of variables that are hypothesized to influence the murder rate.⁴ The data to be examined are state-by-state murder rates in 1950. The variables are divided into three sets. There are four deterrent variables that characterize the criminal justice system, or in economic parlance, the expected out-of-pocket cost of crime. There are four economic variables that measure the opportunity cost of crime. And there are four social/environmental variables that possibly condition the taste for crime. This leaves unmeasured only the expected rewards for criminal behavior, though these are possibly related to the economic and social variables and are otherwise assumed not to vary from state to state.

A simple regression of the murder rate on all these variables leads to the conclusion that each additional execution deters thirteen murders, with a standard error of seven. That seems like such a healthy rate of return, we might want just to randomly draft executees from the population at large. This proposal would be unlikely to withstand the scrutiny of any macroeconomists who are skilled at finding rational expectations equilibria.

The issue I would like to address instead is whether this conclusion is fragile or not. Does it hold up if the list of variables in the model is changed? Individuals with different experiences and different training will find

TABLE 1—VARIABLES USED IN THE ANALYSIS

-
- | | |
|---|--|
| a. Dependent Variable | M = Murder rate per 100,000, FBI estimate. |
| b. Independent Deterrent Variables | PC = (Conditional) Probability of conviction for murder given commission. Defined by $PC = C/Q$, where C = convictions for murder, $Q = M \cdot NS$, NS = state population. This is to correct for the fact that M is an estimate based on a sample from each state.
PX = (Conditional) Probability of execution given conviction (average number of executions 1946–50 divided by C).
T = Median time served in months for murder by prisoners released in 1951.
$XPOS$ = A dummy equal to 1 if $PX > 0$. |
| c. Independent Economic Variables | W = Median income of families in 1949.
X = Percent of families in 1949 with less than one-half W .
U = Unemployment rate.
LF = Labor force participation rate. |
| d. Independent Social and Environmental Variables | NW = Percent nonwhite.
AGE = Percent 15–24 years old.
URB = Percent urban.
$MALE$ = Percent male.
$FAMHO$ = Percent of families that are husband and wife both present families.
$SOUTH$ = A dummy equal to 1 for southern states (Alabama, Arkansas, Delaware, Florida, Kentucky, Louisiana, Maryland, Mississippi, North Carolina, Oklahoma, South Carolina, Tennessee, Texas, Virginia, West Virginia). |
| e. Weighting Variable | $SQRTNF$ = Square root of the population of the FBI-reporting region. Note that weighting is done by multiplying variables by $SQRTNF$. |
| f. Level of Observation | Observations are for 44 states, 35 executing and 9 nonexecuting. The executing states are: Alabama, Arizona, Arkansas, California, Colorado, Connecticut, Delaware, Florida, Illinois, Indiana, Kansas, Kentucky, Louisiana, Maryland, Massachusetts, Mississippi, Missouri, Nebraska, Nevada, New Jersey, New Mexico, New York, North Carolina, Ohio, Oklahoma, Oregon, Pennsylvania, South Carolina, South Dakota, Tennessee, Texas, Virginia, Washington, West Virginia.
The nonexecuting states are: Idaho, Maine, Minnesota, Montana, New Hampshire, Rhode Island, Utah, Wisconsin, Wyoming. |
-

⁴This material is taken from a study by a student of mine, Walter McManus (1982).

different subsets of the variables to be candidates for omission from the equation. Five different lists of doubtful variables are reported in Table 2. A right winger expects

TABLE 2—ALTERNATIVE PRIOR SPECIFICATIONS

Prior	PC	PX	T	XPOS	W	X	U	LF	NW	AGE	URB	MALE	FAMHO	SOUTH
Right Winger	I	I	I	*	D	D	D	D	D	D	D	D	D	D
Rational Maximizer	I	I	I	*	I	I	I	I	D	D	D	D	D	D
Eye-for-an-Eye	I	I	D	*	D	D	D	D	D	D	D	D	D	D
Bleeding Heart	D	D	D	*	I	I	I	I	D	D	D	D	D	D
Crime of Passion	D	D	D	*	I	I	I	I	I	I	I	I	I	I

Notes: 1) I indicates variables considered important by a researcher with the respective prior. Thus, every model considered by the researcher will include these variables. D indicates variables considered doubtful by the researcher. * indicates XPOS, the dummy equal to 1 for executing states. Each prior was pooled with the data two ways: one with XPOS treated as important, and one with it as doubtful.

2) With five basic priors and XPOS treated as doubtful or important by each, we get ten alternative prior specifications.

the punishment variables to have an effect, but treats all other variables as doubtful. He wants to know whether the data still favor the large deterrent effect, if he omits some of these doubtful variables. The rational maximizer takes the variables that measure the expected economic return of crime as important, but treats the taste variables as doubtful. The eye-for-an-eye prior treats all variables as doubtful except the probability of execution. An individual with the bleeding heart prior sees murder as the result of economic impoverishment. Finally, if murder is thought to be a crime of passion then the punishment variables are doubtful.

In Table 3, I have listed the extreme estimates that could be found by each of these groups of researchers. The right-winger minimum of -22.56 means that a regression of the murder rate data on the three punishment variables and a suitably selected linear combination of the other variables yields an estimate of the deterrent effect equal to 22.56 lives per execution. It is possible also to find an estimate of -.86. Anything between these two extremes can be similarly obtained; but no estimate outside this interval can be generated no matter how the doubtful variables are manipulated (linearly). Thus the right winger can report that the inference from this data set that executions deter murders is not fragile. The rational maximizer similarly finds that conclusion insensitive to choice of model, but the other three priors allow execution actually to encourage murder, possibly by a brutalizing effect on society.

TABLE 3—EXTREME ESTIMATES OF THE EFFECT OF EXECUTIONS ON MURDERS

Prior	Minimum Estimate	Maximum Estimate
Right Winger	-22.56	-.86
Rational Maximizer	-15.91	-10.24
Eye-for-an-Eye	-28.66	1.91
Bleeding Heart	-25.59	12.37
Crime of Passion	-17.32	4.10

Note: Least squares is -13.22 with a standard error of 7.2.

I come away from a study of Table 3 with the feeling that any inference from these data about the deterrent effect of capital punishment is too fragile to be believed. It is possible credibly to narrow the set of assumptions, but I do not think that a credibly large set of alternative assumptions will lead to a sharp set of estimates. In another paper (1982), I found a narrower set of priors still leads to inconclusive inferences. And I have ignored the important simultaneity issue (the death penalty may have been imposed in crime ridden states to deter murder) which is often a source of great inferential fragility.

VI. Conclusions

After three decades of churning out estimates, the econometrics club finds itself under critical scrutiny and faces incredulity as never before. Fischer Black writes of "The Trouble with Econometric Models." David

Hendry queries "Econometrics: Alchemy or Science?" John W. Pratt and Robert Schlaifer question our understanding of "The Nature and Discovery of Structure." And Christopher Sims suggests blending "Macroeconomics and Reality."

It is apparent that I too am troubled by the fumes which leak from our computing centers. I believe serious attention to two words would sweeten the atmosphere of econometric discourse. These are whimsy and fragility. In order to draw inferences from data as described by econometric texts, it is necessary to make whimsical assumptions. The professional audience consequently and properly withholds belief until an inference is shown to be adequately insensitive to the choice of assumptions. The haphazard way we individually and collectively study the fragility of inferences leaves most of us unconvinced that any inference is believable. If we are to make effective use of our scarce data resource, it is therefore important that we study fragility in a much more systematic way. If it turns out that almost all inferences from economic data are fragile, I suppose we shall have to revert to our old methods lest we lose our customers in government, business, and on the boardwalk at Atlantic City.

REFERENCES

- Bernard, C., *An Introduction to the Study of Experimental Method*, New York: Mac-Millan, 1927.
- Black, Fischer, "The Trouble with Econometric Models," *Financial Analysts Journal*, March/April 1982, 35, 3-11.
- Friedman, Milton, *A Theory of the Consumption Function*, Princeton: Princeton University Press, 1957.
- Hendry, David, "Econometrics—Alchemy or Science?," *Economica*, November 1980, 47, 387-406.
- Hoerl, Arthur E. and Kennard, Robert W., "Ridge Regression: Biased Estimation for Nonorthogonal Problems," *Technometrics*, February 1970, 12, 55-67.
- Kuhn, Thomas S., *The Structure of Scientific Revolutions*, Chicago: University of Chicago Press, 1962.
- Lakatos, Imre, "Falsification and the Methodology of Scientific Research Programmes," in his and A. Musgrave, eds., *Criticism and the Growth of Knowledge*, Cambridge: Cambridge University Press, 1969.
- Leamer, Edward E., "A Class of Prior Distributions and Distributed Lag Analysis," *Econometrica*, November 1972, 40, 1059-81.
- , "False Models and Post-data Model Construction," *Journal American Statistical Association*, March 1974, 69, 122-31.
- , *Specification Searches: Ad Hoc Inference with Non-experimental Data*, New York: Wiley, 1978.
- , "Techniques for Estimation with Incomplete Assumptions," *IEEE Conference on Decision and Control*, San Diego, December 1981.
- , "Sets of Posterior Means with Bounded Variance Priors," *Econometrica*, May 1982, 50, 725-36.
- McManus, Walter, "Bayesian Estimation of the Deterrent Effect of Capital Punishment," mimeo., University of California-Los Angeles, 1981.
- Plott, Charles R. and Smith, Vernon L., "An Experimental Examination of Two Exchange Institutions," *Review of Economic Studies*, February 1978, 45, 133-53.
- Polanyi, Michael, *Personal Knowledge*, New York: Harper and Row, 1964.
- Pratt, John W. and Schlaifer, Robert, "On the Nature and Discovery of Structure," mimeo., 1979.
- Shiller, Robert, "A Distributed Lag Estimator Derived From Smoothness Priors," *Econometrica*, July 1973, 41, 775-88.
- Sims, C. A., "Macroeconomics and Reality," *Econometrica*, January 1980, 48, 1-48.
- , "Scientific Standards in Econometric Modeling," mimeo., 1982.
- Smith, Vernon L., "Relevance of Laboratory Experiments to Testing Resource Allocation Theory," in J. Kmenta and J. Ramsey, eds., *Evaluation of Econometric Models*, New York: Academic Press, 1980, 345-77.

Exhibit H

Tantalus on the Road to Asymptopia

Edward E. Leamer

My first reaction to “The Credibility Revolution in Empirical Economics,” authored by Joshua D. Angrist and Jörn-Steffen Pischke, was: Wow! This paper makes a stunningly good case for relying on purposefully randomized or accidentally randomized experiments to relieve the doubts that afflict inferences from nonexperimental data. On further reflection, I realized that I may have been overcome with irrational exuberance. Moreover, with this great honor bestowed on my “con” article, I couldn’t easily throw this child of mine overboard.

We economists trudge relentlessly toward Asymptopia, where data are unlimited and estimates are consistent, where the laws of large numbers apply perfectly and where the full intricacies of the economy are completely revealed. But it’s a frustrating journey, since, no matter how far we travel, Asymptopia remains infinitely far away. Worst of all, when we feel pumped up with our progress, a tectonic shift can occur, like the Panic of 2008, making it seem as though our long journey has left us disappointingly close to the State of Complete Ignorance whence we began.

The pointlessness of much of our daily activity makes us receptive when the Priests of our tribe ring the bells and announce a shortened path to Asymptopia. (Remember the Cowles Foundation offering asymptotic properties of simultaneous equations estimates and structural parameters?) We may listen, but we don’t hear, when the Priests warn that the new direction is only for those with Faith, those with complete belief in the Assumptions of the Path. It often takes years down the Path, but sooner or later, someone articulates the concerns that gnaw away in each of

■ *Edward E. Leamer is Professor of Economics, Management and Statistics, University of California at Los Angeles, Los Angeles, California. His e-mail address is (edward.leamer@anderson.ucla.edu).*

doi=10.1257/jep.24.2.31

us and asks if the Assumptions are valid. (T. C. Liu (1960) and Christopher Sims (1980) were the ones who proclaimed that the Cowles Emperor had no clothes.) Small seeds of doubt in each of us inevitably turn to despair and we abandon that direction and seek another.

Two of the latest products-to-end-all-suffering are nonparametric estimation and consistent standard errors, which promise results without assumptions, as if we were already in Asymptopia where data are so plentiful that no assumptions are needed. But like procedures that rely explicitly on assumptions, these new methods work well in the circumstances in which explicit or hidden assumptions hold tolerably well and poorly otherwise. By disguising the assumptions on which nonparametric methods and consistent standard errors rely, the purveyors of these methods have made it impossible to have an intelligible conversation about the circumstances in which their gimmicks do not work well and ought not to be used. As for me, I prefer to carry parameters on my journey so I know where I am and where I am going, not travel stoned on the latest euphoria drug.

This is a story of Tantalus, grasping for knowledge that remains always beyond reach. In Greek mythology Tantalus was favored among all mortals by being asked to dine with the gods. But he misbehaved—some say by trying to take divine food back to the mortals, some say by inviting the gods to a dinner for which Tantalus boiled his son and served him as the main dish. Whatever the etiquette faux pas, Tantalus was punished by being immersed up to his neck in water. When he bowed his head to drink, the water drained away, and when he stretched up to eat the fruit hanging above him, wind would blow it out of reach. It would be much healthier for all of us if we could accept our fate, recognize that perfect knowledge will be forever beyond our reach and find happiness with what we have. If we stopped grasping for the apple of Asymptopia, we would discover that our pool of Tantalus is full of small but enjoyable insights and wisdom.

Can we economists agree that it is extremely hard work to squeeze truths from our data sets and what we genuinely understand will remain uncomfortably limited? We need words in our methodological vocabulary to express the limits. We need sensitivity analyses to make those limits transparent. Those who think otherwise should be required to wear a scarlet-letter *O* around their necks, for “overconfidence.” Angrist and Pischke obviously know this. Their paper is peppered with concerns about quasi-experiments and with criticisms of instrumental variables thoughtlessly chosen. I think we would make progress if we stopped using the words “instrumental variables” and used instead “surrogates”—meaning surrogates for the experiment that we wish we could have conducted. The psychological power of the vocabulary requires a “surrogate” to be chosen with much greater care than an “instrument.”

As Angrist and Pischke persuasively argue, either purposefully randomized experiments or accidentally randomized “natural” experiments can be extremely helpful, but Angrist and Pischke seem to me to overstate the potential benefits of the approach. Since hard and inconclusive thought is needed to transfer the

results learned from randomized experiments into other domains, there must therefore remain uncertainty and ambiguity about the breadth of application of any findings from randomized experiments. For example, how does Card's (1990) study of the effect on the Miami labor market of the Mariel boatlift of 125,000 Cuban refugees in 1980 inform us of the effects of a 2000 mile wall along the southern border of the United States? Thoughts are also needed to justify the choice of instrumental variables, and a critical element of doubt and ambiguity necessarily afflicts any instrumental variables estimate. (You and I know that truly consistent estimators are imagined, not real.) Angrist and Pischke understand this. But their students and their students' students may come to think that it is enough to wave a clove of garlic and chant "randomization" to solve all our problems just as an earlier cohort of econometricians have acted as if it were enough to chant "instrumental variable."

I will begin this comment with some thoughts about the inevitable limits of randomization, and the need for sensitivity analysis in this area, as in all areas of applied empirical work. To be provocative, I will argue here that the financial catastrophe that we have just experienced powerfully illustrates a reason why extrapolating from natural experiments will inevitably be hazardous. The misinterpretation of historical data that led rating agencies, investors, and even myself to guess that home prices would decline very little and default rates would be tolerable even in a severe recession should serve as a caution for all applied econometrics. I will also offer some thoughts about how the difficulties of applied econometric work cannot be evaded with econometric innovations, offering some under-recognized difficulties with instrumental variables and robust standard errors as examples. I conclude with some comments about the shortcomings of an experimentalist paradigm as applied to macroeconomics, and with some warnings about the willingness of applied economists to apply push-button methodologies without sufficient hard thought regarding their applicability and shortcomings.

Randomization Is Not Enough

Angrist and Pischke offer a compelling argument that randomization is one large step in the right direction. Which it is! But like all the other large steps we have already taken, this one doesn't get us where we want to be.

In addition to randomized treatments, most scientific experiments also have controls over the important confounding effects. These controls are needed to improve the accuracy of the estimate of the treatment effect and also to determine clearly the range of circumstances over which the estimate applies. (In a laboratory vacuum, we would find that a feather falls as fast as a bowling ball. In the real world with air, wind, and humidity, all bets are off, pending further study.)

In place of experimental controls, economists can, should, and usually do include control variables in their estimated equations, whether the data are

nonexperimental or experimental. To make my point about the effect of these controls it will be helpful to refer to the prototypical model:

$$y_t = \alpha + (\beta_0 + \beta_1' \mathbf{z}_t) x_t + \boldsymbol{\theta}' \mathbf{w}_t + \varepsilon_t,$$

where x is the treatment, y the response, \mathbf{z} is a set of interactive confounders, \mathbf{w} is a set of additive confounders, where ε stands for all the other unnamed, unmeasured effects which we sheepishly assume behaves like a random variable, distributed independently of the observables. Here $(\beta_0 + \beta_1' \mathbf{z}_t)$ is the variable treatment effect that we wish to estimate. One set of problems is caused by the additive confounding variables \mathbf{w} , which can be uncomfortably numerous. Another set of problems is caused by the interactive confounding variables \mathbf{z} , which may include features of the experimental design as well as characteristics of the subjects.

Consider first the problem of the additive confounders. We have been taught that experimental randomization of the treatment eliminates the requirement to include additive controls in the equation because the correlation between the controls and the treatment is zero by design and regression estimates with or without the controls are unbiased, indeed identical. That's true in Asymptopia, but it's not true here in the Land of the Finite Sample where correlation is an ever-present fact of life and where issues of sensitivity of conclusions to assumptions can arise even with randomized treatments if the correlations between the randomized treatment and the additive confounders, by chance, are high enough.

Indeed, if the number of additive confounding variables is equal to or larger than the number of observations, any treatment x , randomized or not, will be perfectly collinear with the confounding variables (the undersized sample problem). Then, to estimate the treatment effect, we would need to make judgments about which of the confounding variables to exclude. That would ordinarily require a sensitivity analysis, unless through Divine revelation economists were told exactly which controls to include and which to exclude. Though the number of randomized trials may be large, an important sensitivity question can still arise because the number of confounding variables can be increased without limit by using lagged values and nonlinear forms. In other words, if you cannot commit to some notion of smoothness of the functional form of the confounders and some notion of limited or smooth time delays in response, you will not be able to estimate the treatment effect even with a randomized experiment, unless experimental controls keep the confounding variables constant or Divine inspiration allows you to omit some of the variables.

You are free to dismiss the preceding paragraph as making a mountain out of a molehill. By reducing the realized correlation between the treatment and the controls, randomization allows a larger set of additive control variables to be included before we confront the sensitivity issues caused by collinearity. For that reason, though correlation between the treatment and the confounders with nonexperimental data is a *huge* problem, it is much less important when the treatment is randomized. With that problem neutralized, concern shifts elsewhere.

The big problem with randomized experiments is not additive confounders; it's the interactive confounders. This is the heterogeneity issue that especially concerns Heckman (1992) and Deaton (2008) who emphasized the need to study "causal mechanisms," which I am summarizing in terms of the interactive \mathbf{z} variables. Angrist and Pischke completely understand this point, but they seem inappropriately dismissive when they accurately explain "extrapolation of causal effects to new settings is always speculative," which is true, but the extrapolation speculation is more transparent and more worrisome in the experimental case than in the nonexperimental case.

After all, in nonexperimental nonrandomized settings, when judicious choice of additive confounders allows one to obtain just about any estimate of the treatment effect, there is little reason to worry about "extrapolation of causal effects to new settings." What's to extrapolate anyway? Our lack of knowledge? Greater concern about extrapolation is thus an indicator of the progress that comes from randomization.

When the randomization is accidental, we may pretend that the instrumental variables estimator is consistent, but we all know that the assumptions that justify that conclusion cannot possibly hold exactly. Those who use instrumental variables would do well to anticipate the inevitable barrage of questions about the appropriateness of their instruments. Ever-present asymptotic bias casts a large dark shadow on instrumental variables estimates and is what limits the applicability of the estimate even to the setting that is observed, not to mention extrapolation to new settings. In addition, small sample bias of instrumental variables estimators, even in the consistent case, is a huge neglected problem with practice, made worse by the existence of multiple weak instruments. This seems to be one of the points of the Angrist and Pischke paper—purposeful randomization is better than accidental randomization.

But when the randomization is purposeful, a whole new set of issues arises—experimental contamination—which is much more serious with human subjects in a social system than with chemicals mixed in beakers or parts assembled into mechanical structures. Anyone who designs an experiment in economics would do well to anticipate the inevitable barrage of questions regarding the valid transference of things learned in the lab (one value of \mathbf{z}) into the real world (a different value of \mathbf{z}).

With interactive confounders explicitly included, the overall treatment effect $\beta_0 + \beta'\mathbf{z}_i$ is not a number but a variable that depends on the confounding effects. Absent observation of the interactive compounding effects \mathbf{z} , what is estimated is some kind of average treatment effect which is called by Imbens and Angrist (1994) a "Local Average Treatment Effect," which is a little like the lawyer who explained that when he was a young man he lost many cases he should have won but as he grew older he won many that he should have lost, so that on the average justice was done. In other words, if you act as if the treatment effect is a random variable by substituting β_i for $\beta_0 + \beta'\mathbf{z}_i$, the notation inappropriately relieves you of the heavy burden of considering what are the interactive confounders and finding some way to measure them. Less elliptically, absent observation of \mathbf{z} , the estimated treatment

effect should be transferred *only* into those settings in which the confounding interactive variables have values close to the mean values in the experiment. If little thought has gone into identifying these possible confounders, it seems probable that little thought will be given to the limited applicability of the results in other settings. This is the error made by the bond rating agencies in the recent financial crash—they transferred findings from one historical experience to a domain in which they no longer applied because, I will suggest, social confounders were not included. More on this below.

Sensitivity Analysis and Sensitivity Conversations are What We Need

I thus stand by the view in my 1983 essay that econometric theory promises more than it can deliver, because it requires a complete commitment to assumptions that are actually only half-heartedly maintained. The only way to create credible inferences with doubtful assumptions is to perform a sensitivity analysis that separates the fragile inferences from the sturdy ones: those that depend substantially on the doubtful assumptions and those that do not. Since I wrote my “con in econometrics” challenge much progress has been made in economic theory and in econometric theory and in experimental design, but there has been little progress technically or procedurally on this subject of sensitivity analyses in econometrics. Most authors still support their conclusions with the results implied by several models, and they leave the rest of us wondering how hard they had to work to find their favorite outcomes and how sure we have to be about the instrumental variables assumptions with accidentally randomized treatments and about the extent of the experimental bias with purposefully randomized treatments. It’s like a court of law in which we hear only the experts on the plaintiff’s side, but are wise enough to know that there are abundant arguments for the defense.

I have been making this point in the econometrics sphere since before I wrote *Specification Searches: Ad Hoc Inference with Nonexperimental Data* in 1978.¹ That book was stimulated by my observation of economists at work who routinely pass their data through the filters of many models and then choose a few results for reporting purposes. The range of models economists are willing to explore creates ambiguity in the inferences that can properly be drawn from our data, and I have been recommending mathematical methods of sensitivity analysis that are intended to determine the limits of that ambiguity.

¹ Parenthetically, if you are alert, you might have been unsettled by the use of the word “with” in my title: *Ad Hoc Inference with Nonexperimental Data*, since inferences are made *with tools* but *from data*. That is my very subtle way of suggesting that knowledge is created by an interactive exploratory process, quite unlike the preprogrammed estimation dictated by traditional econometric theory.

The language that I used to make the case for sensitivity analysis seems not to have penetrated the consciousness of economists. What I called “extreme bounds analysis” in my 1983 essay is a simple example that is the best-known approach, though poorly understood and inappropriately applied. Extreme bounds analysis is not an “ad hoc but intuitive approach,” as described by Angrist and Pischke. It is a solution to a clearly and precisely defined sensitivity question, which is to determine the range of estimates that the data could support given a precisely defined range of assumptions about the prior distribution. It’s a correspondence between the assumption space and the estimation space. Incidentally, if you could see the wisdom in finding the range of estimates that the data allow, I would work to provide tools that identify the range of *t*-values, a more important measure of the fragility of the inferences.

A prior distribution is a foreign concept for most economists, and I tried to create a bridge between the logic of the analysis and its application by expressing the bounds in language that most economists could understand. Here it is: Include in the equation the treatment variable and a single linear combination of the additive controls. Then find the linear combination of controls that provides the greatest estimated treatment effect and the linear combination that provides the smallest estimated treatment effect. That corresponds to the range of estimates that can be obtained when it is known that the controls are doubtful (zero being the most likely estimate, *a priori*) but there is complete ambiguity about the probable importance of the variables, arbitrary scales, and arbitrary coordinate systems. What is ad hoc are the follow-on methods, for example, computing the standard errors of the bounds or reporting a distribution of estimates as in Sala-i-Martin (1997).

A culture that insists on statistically significant estimates is not naturally receptive to another reason our data are uninformative (too much dependence on arbitrary assumptions). One reason these methods are rarely used is their honesty seems destructive; or, to put it another way, a fanatical commitment to fanciful formal models is often needed to create the appearance of progress. But we need to change the culture and regard the finding of “no persuasive evidence in these data” on the same footing as a “statistically significant and sturdy estimate.” Keep in mind that the sensitivity correspondence between assumptions and inferences can go in either direction. We can ask what set of inferences corresponds to a particular set of assumptions, but we can also ask what assumptions are needed to support a hoped-for inference. That is exactly what an economic theorem does. The intellectual value of the Factor Price Equalization Theorem does not derive from its truthfulness; its value comes from the fact that it provides a minimal set of assumptions that imply Factor Price Equalization, thus focusing attention on *why* factor prices are not equalized.

To those of you who do data analysis, I thus pose two questions that I think every empirical enterprise should be able to answer: What feature of the data leads to that conclusion? What set of assumptions is essential to support that inference?

The Troubles on Wall Street and Three-Valued Logic

With the ashes of the mathematical models used to rate mortgage-backed securities still smoldering on Wall Street, now is an ideal time to revisit the sensitivity issues. Justin Fox (2009) suggests in his title *The Myth of the Rational Market, A History of Risk, Reward, and Delusion on Wall Street* that we have been making a modeling error and that the problems lie in the assumption of rational actors, which presumably can be remedied by business-as-usual after adding a “behavioral” variable or two into the model. I think the roots of the problem are deeper, calling for a change in the way we do business and calling for a book that might be titled: *The Myth of the Data Generating Process: A History of Delusion in Academia*. Rationality of financial markets is a pretty straightforward consequence of the assumption that financial returns are drawn from a “data generating process” whose properties are apparent to experienced investors and econometricians, after studying the historical data. If we don’t know the data generating process, then the efficient markets edifice falls apart. Even simple-minded finance ideas, like the benefits from diversification, become suspect if we cannot reliably assess predictive means, variances, and covariances.

But it isn’t just finance that rests on this myth of a data-generating process. It is the whole edifice of empirical economics. Let’s face it. The evolving, innovating, self-organizing, self-healing human system we call the economy is not well described by a fictional “data-generating process.” The point of the sensitivity analyses that I have been advocating begins with the admission that the historical data are compatible with countless alternative data-generating models. If there is one, the best we can do is to get close; we are never going to know it.

Confronted with our collective colossal failure to anticipate the problems with mortgage-backed securities, we economists have been stampeding shamelessly back to Keynesian thinking about macroeconomics, scurrying to reread Keynes’ (1936) *General Theory of Employment, Interest and Money*. We would do well to go back a little further in time to 1921 when both Keynes’ *Treatise on Probability* and Frank Knight’s *Risk, Uncertainty and Profit* were published. Both of these books are about the myth of the data generating process. Both deal with the limits of the expected utility maximization paradigm. Both serve as foundations for the arguments in favor of sensitivity analysis in my “con” paper. Both are about “three-valued logic.”

Here is how three-valued logic works. Suppose you can confidently determine that it is a good idea to bring your umbrella if the chance of rain is 10 percent or higher, but carrying the umbrella involves more cost than expected benefit if the chance of rain is less. When the data point clearly to a probability either in excess of 10 percent or less than 10 percent, then we are in a world of Knightian risk in which the decision can be based on expected utility maximization. But suppose there is one model that suggests the probability is 15 percent while another equally good model suggests the probability is 5 percent. Then we are in a world of Knightian

uncertainty in which expected utility maximization doesn't produce a decision. When any number between 0.05 and 0.15 is an equally good² assessment of the chance of rain we are dealing with epistemic probabilities that are not numbers but intervals, in the spirit of Keynes (1921). While the decision is two-valued—you either take your umbrella or you don't—the state of mind is three-valued: yes (take the umbrella), no (leave it behind), or I don't know. The point of adopting three-valued logic like Keynes and Knight is to encourage us to think clearly about the limits of our knowledge and the limits of the expected utility maximization paradigm. With all of the focus on decision making with two-valued logic, the profession has done precious little work on decision under ambiguity and three-valued logic.³

In other words, when I asked us to “take the con out of econometrics” I was only saying the obvious: If the range of inferences that can reasonably be supported by the data we have is too wide to point to one and only one decision, we need to admit that the data leave us confused. Thus my contribution to econometrics has been confusion! You, though, refuse to admit that you are just as confused as I.

Three-valued logic seems especially pertinent when extending the results of experiments into other domains—there is a lot of “we don't know” there. Even in the case of mechanical systems, with statistical properties much better understood than economics/financial human systems, it is not assumed that models will work under extreme conditions. Engineers may design aircraft that according to their computer models can fly, but until real airplanes are actually tested in normal and stressful conditions, the aircraft are not certified to carry passengers. Indeed, Boeing's composite plastic 787 Dreamliner has been suffering production delays because wing damage has shown up “when the stress on the wings was well below the load the wings must bear to be federally certified to carry passengers” (Gates, 2009). Too bad we couldn't have stress-tested those mortgage-backed securities before we started flying around the world with them. Too bad the rating agencies did not use three-valued logic with another bond rating: “AAA-H,” meaning hypothetically AAA according to a model but not yet certified as recession-proof.


There is of course more than one reason why mortgage-backed securities didn't fly when the weather got rough, but it will suit my purposes here if I make the argument that it was an inappropriate extrapolation of data from one accidental experiment to a different setting that is at the root of the problem. To make this argument, it is enough to look at the data in Los Angeles. Don't expect a full econometric housing model. It's only a provocative illustration.

Figure 1 contrasts nonfarm payrolls in Los Angeles and in the United States overall during the last three recessions. In 2001 and in 2008, the L.A. job market

² Please don't think I am assigning a uniform distribution over this interval, since if that were the case, the probability would be precisely equal to the mean: 0.10.

³ Bewley (1986), Klibanoff, Marinacci, and Mukerji (2005), and Hanany and Klibanoff (2009) and references therein are exceptions.

Figure 1

Payroll Jobs in Los Angeles and U.S. Overall*(recessions shaded)*


Note: The figure shows the change in number of employees on nonfarm payrolls, in the Los Angeles area and in the U.S. overall, relative to January 1990, seasonally adjusted. The Los Angeles area is the Los Angeles–Long Beach–Santa Ana, California, Metropolitan Statistical Area.

declined in parallel with the U.S. decline, but the recession of 1990–91 was especially severe in Los Angeles. U.S. payroll jobs dropped by 1.5 percent during this recession, while jobs in the Los Angeles metropolitan statistical area declined by 11 percent and did not return to their 1990 levels until 1999. This was a natural experiment known as the end of the Cold War, with Los Angeles treated and with, for example, San Francisco in the control group. (The 2001 recession had the treatment reversed, with San Francisco treated to a tech bust but Los Angeles in the control group.)

Thus, I suggest, the L.A. data in the early 1990s is a test case of the effect of a severe recession on mortgage defaults. Figure 2 illustrates the number of homes sold and the median prices in Los Angeles from January 1985 to August 2009. The horizontal line segments indicate the period over which all homes purchased with a loan to value ratio of 90 percent (indicating a 10 percent downpayment) and interest-only payments will be underwater at the next trough. If all these homes were returned to the banks under the worst case scenario when the price hits bottom, the bank losses are 90 percent of the gap between that line segment and the price at origination. Clearly the underwater problem is much more intense in the latest downturn than it was in the 1990s.

In the first episode, volume peaked much before prices, falling by 50 percent in 28 months. Though volume was crashing, the median price continued to increase, peaking in May 1991, 101 percent above its value at the start of these data in January 1985. Then commenced a slow price decline with a trough in

Figure 2

Los Angeles Housing Market: Median Home Price and Sales Volume

Line segments identify underwater homes at cycle troughs assuming 90% loan to value

Source: California Association of Realtors.

Note: The median price and sales volume are a 3-month moving average, seasonally adjusted. The line segments identify homes that will be underwater at the coming cycle trough assuming 90 percent loan to value at origination and interest only payments. (Denoting the price at time t by p_t and the price at the trough by p_{\min} , with 10 percent down, the loan balance is $0.9 p_t$, and the home is underwater at the trough if $0.9 p_t > p_{\min}$ or if $p_t > p_{\min}/0.9$.)


December 1996, with the median price 29 percent below its previous peak, a decline of 5 percent per year.

I have been fond of summarizing these data by saying that for homes, it's a volume cycle, not a price cycle. This very slow price-discovery occurs because people celebrate investment gains, but deny losses. Owner-occupants of homes can likewise hold onto long-ago valuations and insist on prices that the market cannot support. Because of that denial, there are many fewer transactions, and the transactions that do occur tend to be at the seller's prices, not equilibrium market prices.

The slow price discovery acts like a time-out, allowing the fundamentals to catch up to valuations and keeping foreclosure rates at minimal levels.⁴ In the early phase of the current housing correction, history seemed to be repeating itself, since volume was dropping rapidly even as prices continued to rise. But then began a rapid 51 percent drop in home prices between August 2007 to March 2009, creating a huge amount of underwater valuation.

⁴ The first underwater problem illustrated in Figure 2 could be almost completely remedied by a switch from 90 to 80 percent loan-to-value ratios starting when volumes began to drop, because most of the underwater loans came after that break in the market.

Figure 3
Southern California Foreclosures



Source: MDA DataQuick.

Note: The concentration of foreclosures in time and space in Southern California in the latest housing correction is illustrated in Figure 3, which displays quarterly foreclosures since 1987 for both Los Angeles County and “The Inland Empire,” composed of the two “peripheral” counties east of Los Angeles County (Riverside and San Bernadino), where subprime lending was very prevalent. In both Los Angeles and the Inland Empire in the 1990s, the rise in foreclosures trailed the price movement by several years. In contrast, the time-concentrated spike in foreclosures in 2007 occurred at the very start of the price erosion and has been much more extreme in the Inland Empire than in Los Angeles.

Why did the L.A. 1990s data mislead with regard to the current housing correction? One possible answer is untested social effects and unmeasured subject effects—two very important interactive confounders. Innovations in mortgage origination in 2003–2005 extended the home-ownership peripheries of our cities both in terms of income and location. When the subprime mortgage window shut down, the demand for owner-occupied homes at the extended peripheries was eliminated virtually overnight. Since these properties were financed with loans that could only work if the houses paid for themselves via appreciation, the banks became the new owners. Accounting rules do not allow banks to deny losses the way owner-occupants do, and banks immediately dumped the foreclosed homes onto the market at the worst time, in the worst way, with broken windows and burned-up lawns, concentrated in time and space, causing very rapid (or even exaggerated) price discovery in the affected peripheries. In contrast, foreclosures in the 1990s, illustrated in Figure 3 for the case of Southern California, were delayed, and were dispersed in time and space.

We all need to learn both narrower and broader lessons here. I expected price discovery in housing markets to be slow, as it had been after the bursting of previous bubbles, and I was completely wrong. I will be more careful about interactive social confounders in the future.

White-Washing

It should not be a surprise at this point in this essay that I part ways with Angrist and Pischke in their apparent endorsement of White's (1980) paper on how to calculate robust standard errors. Angrist and Pischke write: "Robust standard errors, automated clustering, and larger samples have also taken the steam out of issues like heteroskedasticity and serial correlation. A legacy of White's (1980) paper on robust standard errors, one of the most highly cited from the period, is the near-death of generalized least squares in cross-sectional applied work."

An earlier generation of econometricians corrected the heteroskedasticity problems with weighted least squares using weights suggested by an explicit heteroskedasticity model. These earlier econometricians understood that reweighting the observations can have dramatic effects on the actual estimates, but they treated the effect on the standard errors as a secondary matter. A "robust standard" error completely turns this around, leaving the estimates the same but changing the size of the confidence interval. Why should one worry about the length of the confidence interval, but not the location? This mistaken advice relies on asymptotic properties of estimators.⁵ I call it "White-washing." Best to remember that no matter how far we travel, we remain always in the Land of the Finite Sample, infinitely far from Asymptopia. Rather than mathematical musings about life in Asymptopia, we should be doing the hard work of modeling the heteroskedasticity and the time dependence to determine if sensible reweighting of the observations materially changes the locations of the estimates of interest as well as the widths of the confidence intervals.

Estimation with instrumental variables is another case of inappropriate reliance on asymptotic properties. In finite samples, these estimators can seriously distort the evidence for the same reason that the ratio of two sample means, \bar{y}/\bar{x} , is a poor summary of the data evidence about the ratio of the means, $E(y)/E(x)$, when the finite sample leaves a "dividing-by-zero-problem" because the denominator \bar{x} is not statistically far from zero. This problem is greatly amplified with multiple weak instruments, a situation that is quite common. It is actually quite feasible from a Bayesian perspective to program an alert into our instrumental variables estimation together with remedies, though this depends on Assumptions. In other words, you have to do some hard thinking to use instrumental variables methods in finite samples.

As the sample size grows, concern should shift from small-sample bias to asymptotic bias caused by the failure of the assumptions needed to make instrumental variables work. Since it is the unnamed, unobserved variables that are the source of the problem, this isn't easy to think about. The percentage bias is small if

⁵ The change in length but not location of a confidence interval is appropriate for one specialized covariance structure.

the variables you have forgotten are unimportant compared with the variables that you have remembered. That's easy to determine, right?

A Word on Macroeconomics

Finally, I think that Angrist and Pischke are way too optimistic about the prospects for an experimental approach to macroeconomics. Our understanding of causal effects in macroeconomics is virtually nil, and will remain so. Don't we know that? Though many members of our profession have jumped up to support the \$787 billion stimulus program in 2009 as if they knew that was an appropriate response to the Panic of 2008, the intellectual basis for that opinion is very thin, especially if you take a close look at how that stimulus bill was written.

The economists who coined the DSGE acronym combined in three terms the things economists least understand: "dynamic," standing for forward-looking decision making; "stochastic," standing for decisions under uncertainty and ambiguity; and "general equilibrium," standing for the social process that coordinates and influences the actions of all the players. I have tried to make this point in the title of my recent book: *Macroeconomic Patterns and Stories* (Leamer, 2009). That's what we do. We seek patterns and tell stories.

Conclusion

Ignorance is a formidable foe, and to have hope of even modest victories, we economists need to use every resource and every weapon we can muster, including thought experiments (theory), and the analysis of data from nonexperiments, accidental experiments, and designed experiments. We should be celebrating the small genuine victories of the economists who use their tools most effectively, and we should dial back our adoration of those who can carry the biggest and brightest and least-understood weapons. We would benefit from some serious humility, and from burning our "Mission Accomplished" banners. It's never gonna happen.

Part of the problem is that we data analysts want it all automated. We want an answer at the push of a button on a keyboard. We know intellectually that thoughtless choice of an instrument can be a severe problem and that summarizing the data with the "consistent" instrumental variables estimate when the instruments are weak is an equally large error.⁶ The substantial literature on estimation with weak instruments has not yet produced a serious practical competitor to the usual

⁶ Bayesians have a straightforward solution in theory to this problem: describe the marginal likelihood function, marginalized with respect to all the parameters except the coefficient being estimated. If the instruments are strong, this marginal likelihood will have its mode near the instrumental variables estimate. If the instruments are weak, the central tendency of the marginal likelihood will lie elsewhere, sometimes near the ordinary least-squares estimate.

instrumental variables estimator. Our keyboards now come with a highly seductive button for instrumental variables estimates. To decide how best to adjust the instrumental variables estimates for small-sample distortions requires some hard thought. To decide how much asymptotic bias afflicts our so-called consistent estimates requires some very hard thought and dozens of alternative buttons. Faced with the choice between thinking long and hard versus pushing the instrumental variables button, the single button is winning by a very large margin.

Let's not add a "randomization" button to our intellectual keyboards, to be pushed without hard reflection and thought.

■ *Comments from Angus Deaton, James Heckman, Guido Imbens, and the editors and referees for the JEP (Timothy Taylor, James Hines, David Autor, and Chad Jones) are gratefully acknowledged, but all errors of fact and opinion, of course, remain my own.*

References

- Bewley, Truman F.** 1986. "Knightian Decision Theory." Parts 1 and 2. Cowles Foundation Discussion Papers No. 807 and 835.
- Card, David.** 1990. "The Impact of the Mariel Boatlift on the Miami Labor Market." *Industrial and Labor Relations Review*, 43(2): 245–57.
- Chamberlain, Gary, and Guido Imbens.** 1996. "Hierarchical Bayes Models with Many Instrumental Variables." NBER Technical Working Paper 204.
- The Economist.** 2009. "Cause and Defect." August 13, 2009.
- Deaton, Angus.** 2008. "Instruments of Development: Randomization in the Tropics, and the Search for the Elusive Keys to Economic Development." The Keynes Lecture, British Academy, October, 2008.
- Fox, Justin.** 2009. *The Myth of the Rational Market. A History of Risk, Reward and Delusion on Wall Street.* New York: HarperCollins.
- Gates, Dominic.** 2009. "Boeing 787 Wing Flaw Extends Inside Plane." *Seattle Times*, July 30. <http://seattletimes.nwsourc.com/html/boeingaerospace/2009565319-boeing30.html>.
- Hanany, Eran, and Peter Klibanoff.** 2009. "Updating Ambiguity Averse Preferences." *The B.E. Journal of Theoretical Economics*, vol. 9, issue 1 (Advances), article 37.
- Heckman, James J.** 1992. "Randomization and Social Program Evaluation." In *Evaluating Welfare and Training Programs*, eds. Charles Manski and Irwin Garfinkel, 201–230. Cambridge, MA: Harvard University Press. (Also available as NBER Technical Working Paper 107.)
- Imbens, Guido W.** 2009. "Better LATE than Nothing: Some Comments on Deaton (2009) and Heckman and Urzua (2009)." NBER Working Paper No. w14896.
- Imbens, Guido W., and Joshua D. Angrist.** 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica*, 62(2): 467–75.
- Keynes, John Maynard.** 1921. *A Treatise on Probability.* London: Macmillan.
- Keynes, John Maynard.** 1936. *The General Theory of Employment, Interest and Money.* Macmillan.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji.** 2005. "A Smooth Model of Decision Making Under Ambiguity." *Econometrica*, 73(6): 1849–1892.
- Knight, Frank.** 1921. *Risk, Uncertainty and Profit.* New York: Houghton Mifflin.
- Leamer, Edward E.** 1978. *Specification Searches: Ad Hoc Inference with Nonexperimental Data.* John Wiley and Sons.
- Leamer, Edward E.** 1983. "Let's Take the Con

Out of Econometrics.” *American Economic Review*, 73(1): 31–43.

Leamer, Edward E. 1985. “Vector Autoregressions for Causal Inference?” In Karl Brunner and Allan H. Meltzer, eds., *Carnegie-Rochester Conference Series on Public Policy*, vol. 22, pp. 255–304.

Leamer, Edward E. 2009. *Macroeconomic Patterns and Stories*. Springer Verlag.

Liu, Ta-Chung. 1960. “Underidentification, Structural Estimation, and Forecasting.”

Econometrica, 28(4): 855–65.

Sala-i-Martin, Xavier. 1997. “I Just Ran Two Million Regressions.” *American Economic Review*, 87(2): 178–83.

Sims, Christopher. 1980. “Macroeconomics and Reality.” *Econometrica*, 48(1): 1–48.

White, Halbert. 1980. “A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity.” *Econometrica*, 48(4): 817–38.